arXiv:1408.5752v2 [cs.LO] 6 Oct 2014

The Complexity of Counting Models of Linear-time Temporal Logic*,**

Hazem Torfah and Martin Zimmermann

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
{torfah, zimmermann}@react.uni-saarland.de

Abstract. We determine the complexity of counting models of bounded size of specifications expressed in Linear-time Temporal Logic. Counting word-models is #P-complete, if the bound is given in unary, and as hard as counting accepting runs of nondeterministic polynomial space Turing machines, if the bound is given in binary. Counting tree-models is as hard as counting accepting runs of nondeterministic exponential time Turing machines, if the bound is given in unary. For a binary encoding of the bound, the problem is at least as hard as counting accepting runs of nondeterministic exponential space Turing machines. On the other hand, it is not harder than counting accepting runs of nondeterministic doubly-exponential time Turing machines.

1 Introduction

Model counting, the problem of computing the *number of models* of a logical formula, generalizes the satisfiability problem and has diverse applications: many probabilistic inference problems, such as Bayesian net reasoning [13], and planning problems, such as computing the robustness of plans in incomplete domains [15], can be formulated as model counting problems of propositional logic. Model counting for Linear-time Temporal Logic (LTL) has been recently introduced in [8]. LTL is the most commonly used specification logic for reactive systems [16] and the standard input language for model checking [2,5] and synthesis tools [3,4,6]. LTL model counting asks for computing the number of transition systems that satisfy a given LTL formula. As such a formula has either zero or infinitely many models one considers models of *bounded* size: for a formula φ and a bound k, the problem is to count the number of models of φ of size k. This is motivated by applications like bounded model checking [2] and bounded synthesis [7], where one looks for short error paths and small implementations, respectively, by iteratively increasing a bound on the size of the model. Just like propositional model counting generalizes satisfiability, by considering two types of bounded models, namely, *word-models* (of length k) and *tree-models* (of

* This work was partially supported by the German Research Foundation (DFG) as part of SFB/TR 14 AVACS and by the Deutsche Telekom Foundation.

** A short version appears in the Proceedings of FSTTCS 2014.

height k), the authors of [8] introduced quantitative extensions of model checking and synthesis.

Word-models are ultimately periodic words of the form $u.v^\omega$ of bounded length $|u.v|$, which are used to model computations of a system. Counting word-models can be used in *model checking* to determine not only the existence of computations that violate the specification, but also the *number* of such *violations*. To this end, one turns the model checking problem into an LTL satisfiability problem by encoding the transition system and the negation of the specification into a single LTL formula. Its models represent erroneous computations of the system, i.e., counting them gives a quantitative notion of satisfaction.

Tree-models are finite trees (of fixed out-degree) of bounded height with back-edges at the leaves, i.e., tree-models can be exponentially-sized in the bound. They are used to describe implementations of the input-output behavior of reactive systems (see, e.g., [7]), namely the edges of a tree-model represent the input behavior of the environment and the nodes represent the corresponding output behavior of the system. In *synthesis*, counting tree-models can be used to determine not only the existence of an implementation that satisfies the specification, but also the *number* of such *implementations*. This number is a helpful metric to understand how much room for implementation choices is left by a given specification, and to estimate the impact of new requirements on the remaining design space.

For *safety* LTL specifications [18], algorithms solving the word- and the tree-model counting problem were presented in [8]. The running time of the algorithms is linear in the bound and doubly-exponential respectively triply-exponential in the length of the formula. The high complexity in the formula is, however, not a major concern in practice, since specifications are typically small while models are large (cf. the state-space explosion problem).

Here, we complement these algorithms by analyzing the computational complexity of the model counting problems for *full* LTL by placing the problems into counting complexity classes. These classes are based on counting accepting runs of nondeterministic Turing machines. In his seminal paper on the complexity of computing the permanent [19], Valiant introduced the class #P of counting problems associated with counting accepting runs of nondeterministic polynomial time Turing machines: a function $f\colon \Sigma^* \to \mathbb{N}$ is in #P if there is a nondeterministic polynomial time Turing machine $\mathcal{M}$ such that $f(w)$ is equal to the number of accepting runs of $\mathcal{M}$ on w. Furthermore, for a class $\mathcal{C}$ of decision problems, he defined[1] $\#_o\mathcal{C}$ to be the class of counting problems induced by counting accepting runs of a nondeterministic polynomial time Turing machine with an oracle from $\mathcal{C}$.

A nondeterministic polynomial time Turing machine $\mathcal{M}$ (with or without oracle) has at most $\mathcal{O}(2^{p(n)})$ different runs on inputs of length n for some polynomial p. This means that there is an exponential upper bound on functions in #P and in $\#_o\mathcal{C}$ for every $\mathcal{C}$. However, an LTL tautology has exponentially many

[1] Valiant originally used the notation $\#\mathcal{C}$, but we added the subscript to distinguish the oracle-based classes from the classes introduced below.

word-models of length k and more than doubly-exponentially many tree-models of height k. This means, that no function in any of the counting classes defined above can capture the counting problems for LTL.

To overcome this, we consider counting classes obtained by lifting the restriction on considering only nondeterministic polynomial time (oracle) machines: a function $f\colon \Sigma^* \to \mathbb{N}$ is in #Pspace, if there is a nondeterministic polynomial *space* Turing machine $\mathcal{M}$ such that $f(w)$ is equal to the number of accepting runs of $\mathcal{M}$ on w. The classes #Exptime, #Expspace, and #2Exptime are defined analogously[2]. Some of these classes appeared in the literature, e.g., #Pspace was shown to be equal to FPspace [11] (if the output is encoded in binary). Also, computing a specific entry of a matrix power A^n is in #Pspace, if A is represented succinctly and n in binary [14], and counting self-avoiding walks in succinctly represented hypercubes is complete for #Exptime [12] under right-bit-shift reductions.

We place the LTL model counting problems in these classes. Unsurprisingly, the encoding of the bound k is crucial: for unary bounds, we show counting word-models to be #P-complete and show counting tree-models to be #Exptime-complete. For binary bounds, the word-model counting problem is #Pspace-complete and counting tree-models is #Expspace-hard and in #2Exptime. The upper bounds hold for full LTL while the formulas for the lower bounds define safety properties (using only the temporal operators next and release). Thus, the lower bounds already hold for the fragment considered in [8].

The algorithms we present to prove the upper bounds are not practical since they are based on guessing a word (tree) and then model checking it. Hence, a deterministic variant of these algorithms would enumerate all words (trees) of length (height) k and then run a model checking algorithm on them. In particular, the running time of the algorithms is exponential (or worse) in the bound k, which is in stark contrast to the practical algorithms [8]. Our lower bounds are reductions from the problem of counting accepting runs of a Turing machine. For the word counting problem, the reductions are slight strengthenings of the reduction proving Pspace-hardness of the LTL model checking problem [17]. However, the reductions in the tree case are more involved (and to the best of our knowledge new), since we have to deal with exponential time respectively exponential space Turing machines. The main technical difficulties are to encode runs of exponential length and with configurations of exponential size into tree-models of "small" LTL formulas and to ensure that there is a one-to-one correspondence between accepting runs and models of the constructed formula.

All missing proofs can be found in the appendix.

2 Preliminaries

We represent models as *labeled transition systems*. For a given finite set Υ of directions and a finite set Σ of labels, a Σ-labeled Υ-transition system is a tuple

[2] Following the "satanic" [9] tradition of naming counting classes, we drop the N (standing for nondeterministic) in the names of the classes, just as it is done for #P.

$\mathcal{S} = (S, s_0, \tau, o)$, consisting of a finite set of states S, an initial state $s_0 \in S$, a transition function $\tau\colon S \times \Upsilon \to S$, and a labeling function $o\colon S \to \Sigma$. A *path* in $\mathcal{S}$ is a sequence $\pi\colon \mathbb{N} \to S \times \Upsilon$ of states and directions that follows the transition function, i.e., for all $i \in \mathbb{N}$ if $\pi(i) = (s_i, e_i)$ and $\pi(i+1) = (s_{i+1}, e_{i+1})$, then $s_{i+1} = \tau(s_i, e_i)$. We call the path initial if it starts with the initial state: $\pi(0) = (s_0, e)$ for some $e \in \Upsilon$.

We use Linear-time Temporal Logic (LTL) [16], with the usual temporal operators Next $\bigcirc$, Until $\mathcal{U}$, Release $\mathcal{V}$, and the derived operators Eventually $\Diamond$ and Globally $\Box$. We use $\bigcirc^i$ to refer to i nested next operators. LTL formulas are defined over a set of atomic propositions $AP = I \cup O$, which is partitioned into a set I of input propositions and a set O of output propositions. We denote the satisfaction of an LTL formula φ by an infinite sequence $\sigma\colon \mathbb{N} \to 2^{AP}$ of valuations of the atomic propositions by $\sigma \models \varphi$. A 2^O-labeled 2^I-transition system $\mathcal{S} = (S, s_0, \tau, o)$ satisfies φ, if for every initial path π the sequence $\sigma_\pi : i \mapsto o(\pi(i))$, where $o(s, e) = (o(s) \cup e)$, satisfies φ. Then $\mathcal{S}$ is a model of φ.

A *k-word-model* of an LTL formula φ over AP is a pair (u, v) of finite words over 2^{AP} such that $|u.v| = k$ and $u.v^\omega \models \varphi$. We call u the prefix and v the period of (u, v). Note that an ultimately periodic word might be induced by more than one k-word-model, i.e., $\{a\}^\omega$ is induced by the 2-word-models $(\{a\}, \{a\})$ and $(\varepsilon, \{a\}\{a\})$.

A *k-tree-model* of an LTL formula φ over $AP = I \cup O$ is a 2^O-labeled 2^I-transition system that forms a tree (whose root is the initial state) of height k with added back-edges from the leaves (for each leaf and direction, there is an edge to a state on the branch leading to the leaf) that satisfies φ. Figure 1 shows a tree-model of height one.

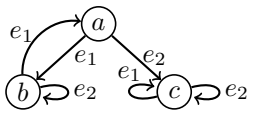

Fig. 1. A tree-model.

Fix $AP = I \cup O$. For a formula φ and $k \in \mathbb{N}$, the *k-word* (*k-tree*) *counting problem* asks to compute the number of k-word-models (k-tree-models up to isomorphism) of φ over AP.

3 Counting Complexity Classes

We use nondeterministic Turing machines with or without oracle access to define counting complexity classes, which we assume (without loss of generality) to terminate on every input. For background on (oracle) Turing machines and counting complexity we refer to [1].

A function $f\colon \Sigma^* \to \mathbb{N}$ is in the class #P [19] if there is a nondeterministic polynomial time Turing machine $\mathcal{M}$ such that $f(w)$ is equal to the number of accepting runs of $\mathcal{M}$ on w. Similarly, for a given complexity class $\mathcal{C}$ of decision problems, a function f is in $\#_o\mathcal{C}$ [19,9] if there is a nondeterministic polynomial time oracle Turing machine $\mathcal{M}$ with oracle in $\mathcal{C}$ such that $f(w)$ is equal to the number of accepting runs of $\mathcal{M}$ on w. As a nondeterministic polynomial time Turing machine $\mathcal{M}$ (with or without oracle) has at most $\mathcal{O}(2^{p(n)})$ runs on inputs of length n for some polynomial p (that only depends on $\mathcal{M}$), we obtain

an exponential upper bound on functions in #P and $\#_o\mathcal{C}$ for every $\mathcal{C}$, which explains the need for larger counting classes to characterize the model counting problems for LTL.

A function $f\colon \Sigma^* \to \mathbb{N}$ is in #PSPACE, if there is a nondeterministic polynomial *space* Turing machine $\mathcal{M}$ such that $f(w)$ is equal to the number of accepting runs of $\mathcal{M}$ on w. #EXPTIME, #EXPSPACE, and #2EXPTIME are defined by counting accepting runs of nondeterministic exponential time, exponential space, and doubly-exponential time machines.

Proposition 1.

1. $\#\mathrm{P} \subseteq \#_o\textsc{Pspace} \subseteq \#_o\textsc{Exptime} \subseteq \#_o\textsc{NExptime} \subseteq \#_o\textsc{Expspace} \subseteq \#_o\textsc{2Exptime}$.
2. $\#\textsc{Pspace} \subseteq \#\textsc{Exptime} \subsetneq \#\textsc{Expspace} \subseteq \#\textsc{2Exptime}$.
3. $f \in \#\textsc{Exptime}$ *implies* $f(w) \in \mathcal{O}(2^{2^{p(|w|)}})$ *for a polynomial* p.
4. $f \in \#\textsc{2Exptime}$ *implies* $f(w) \in \mathcal{O}(2^{2^{2^{p(|w|)}}})$ *for a polynomial* p.
5. $w \mapsto 2^{2^{|w|}}$ *is in* #PSPACE
6. $w \mapsto 2^{2^{2^{|w|}}}$ *is in* #EXPSPACE.

We continue by relating the oracle-based and the generalized classes introduced above.

Lemma 2. $\#_oC \subsetneq \#C$ *for* $C \in \{\textsc{Pspace}, \textsc{Exptime}, \textsc{Expspace}, \textsc{2Exptime}\}$.

Proof. We show $\#_o\textsc{Pspace} \subsetneq \#\textsc{Pspace}$, the other claims are proven analogously. Let $f \in \#_o\textsc{Pspace}$, i.e., there is a nondeterministic polynomial time Turing machine $\mathcal{M}$ with oracle $A \in \textsc{Pspace}$ such that $f(w)$ is equal to the number of accepting runs of $\mathcal{M}$ on w. Note that all oracle queries are polynomially-sized in the length $|w|$ of the input to $\mathcal{M}$, since $\mathcal{M}$ is polynomially time-bounded. Hence, in nondeterministic polynomial space one can simulate $\mathcal{M}$ and evaluate the oracle calls explicitly by running a deterministic machine deciding A in polynomial space. Since the oracle queries are evaluated deterministically, the simulation has as many accepting runs as $\mathcal{M}$ has. Thus, $f \in \#\textsc{Pspace}$.

Now, consider the function $|w| \mapsto 2^{2^{|w|}}$, which is in #PSPACE, but not in $\#_o\textsc{Pspace}$.

We use parsimonious reductions to define hardness and completeness, i.e., the most restrictive notion of reduction for counting problems. A counting problem f is #P-hard, if for every $f' \in \#\mathrm{P}$ there is a polynomial time computable function r such that $f'(x) = f(r(x))$ for all inputs x. In particular, if f' is induced by counting the accepting runs of $\mathcal{M}$, then r depends on $\mathcal{M}$ (and possibly on its time-bound $p(n)$). Furthermore, f is #P-complete, if f is #P-hard and $f \in \#\mathrm{P}$. Hardness and completeness for the other classes are defined analogously.

4 Counting Word-Models

In this section, we provide matching lower and upper bounds for the complexity of counting k-word-models of an LTL specification.

Our hardness proofs are based on constructing an LTL formula $\varphi^w_{\mathcal{M}}$ for a given Turing machine $\mathcal{M}$ and an input w that encodes the accepting runs of $\mathcal{M}$ on w. Constructing such an LTL formula is straightforward and can be done in polynomial time for Turing machines with polynomially-sized configurations [17]. However, the challenge is to construct $\varphi^w_{\mathcal{M}}$ such that the number of accepting runs on w is equal to the number of k-word-models of $\varphi^w_{\mathcal{M}}$ for a fixed bound k. To this end, we have to enforce that each accepting run is represented by a unique k-word-model, i.e., by a unique prefix and period of total length k. We choose k such that a run on w of maximal length can be encoded in $k-1$ symbols and define $\varphi^w_{\mathcal{M}}$ such that it has only k-word-models whose period has length one. If a run of $\mathcal{M}$ is shorter than the maximal-length run we repeat the final configuration until reaching the maximal length, which is achieved by accompanying the configurations in the encoding with consecutive id's.

For the upper bounds we show that there are appropriate nondeterministic Turing machines that guess an ultimately-periodic word and model check it against φ, i.e., the number of accepting runs on k and φ is equal to the number of k-word-models of φ.

4.1 The Case of Unary Encodings.

We show that counting word-models for unary bounds is #P-complete.

Theorem 3. *The following problem is #P-complete: Given an LTL formula φ and a bound k (in unary), how many k-word-models does φ have?*

Proof. We start with the hardness proof. Let $\mathcal{M} = (Q, q_\iota, Q_F, \Sigma, \delta)$ be a one-tape nondeterministic polynomial time Turing machine, where Q is the set of states, q_ι is the initial state, Q_F is the set of accepting states, Σ is the alphabet, and $\delta\colon (Q \setminus Q_F) \times \Sigma \to 2^{Q \times \Sigma \times \{-1,1\}}$ is the transition function, where -1 and 1 encode the directions of the head. Note that the accepting states are terminal and that $\mathcal{M}$ rejects by terminating in a nonaccepting state. Let $\mathcal{M}$ be $p(n)$-time bounded for some polynomial p, and let $w = w_0 \cdots w_{n-1}$ be an input to $\mathcal{M}$. We construct an LTL formula $\varphi^w_{\mathcal{M}}$ and define a bound k, both polynomial in $|w|$ and $|\mathcal{M}|$, such that the number of accepting runs of $\mathcal{M}$ on w is equal to the number of k-word-models of $\varphi^w_{\mathcal{M}}$.

A run of $\mathcal{M}$ on w is encoded by a finite alternating sequence of id's id_i and configurations c_i that is followed by an infinite repetition of a dummy symbol:

$$\$\ \mathrm{id}_0\ \#\ c_0\ \$\ \mathrm{id}_1\ \#\ c_1\ \$\ \mathrm{id}_2\ \#\ c_2\ \$\ \cdots\ \$\ \mathrm{id}_{p(n)}\ \#\ c_{p(n)}\ (\bot)^\omega \tag{1}$$

Note that the period of the word-model is of the form $\bot^\ell$ for some $\ell > 0$. We will define k such that maximal-length runs of $\mathcal{M}$ on w can be encoded in the prefix, and such that the only possible period has length one by ensuring that

exactly $p(n)$ configurations are encoded (by repeating the final configuration if necessary). This ensures that an accepting run is encoded by exactly one k-word-model.

Let $l_r = p(n)$ be the maximal length of a run of $\mathcal{M}$ on w. The size of a configuration of $\mathcal{M}$ on w is also bounded by l_r. For the id's we use an encoding of a binary counter with $l_c = \lceil \log l_r \rceil$ many bits. Let $AP = (Q \cup \Sigma) \cup \{b_1, \ldots, b_{l_c}, \$, \#, \bot\}$ be the set of atomic propositions. The atomic propositions in $Q \cup \Sigma$ are used to encode the configuration of $\mathcal{M}$ by encoding the tape contents, the state of the machine, and the head position. The atomic propositions $b_1, \ldots, b_{l_c}$ represent the bit values of an id. The symbols \$ and # are used as separators between id's and configurations, and $\bot$ is a dummy symbol for the model's period. The distance between two \$ symbols and also between two # symbols in the encoding is given by $d = l_r + 3$ (see (1)). Then, $\varphi^w_{\mathcal{M}}$ is the conjunction of the following formulas:

- *Id* encodes the id's of the configurations. It uses a formula $Inc(b_1, \ldots, b_{l_c}, d)$ that asserts that the number encoded by the bits b_j after d steps is obtained by incrementing the number encoded at the current position. This formula will be reused in the tree case.
- *Init* asserts that the run of $\mathcal{M}$ starts with the initial configuration.
- *Accept* asserts that the run must reach an accepting configuration.
- *Config* declares the consistency of two successive configurations with the transition relation of $\mathcal{M}$. Here, we use d many next operators to relate the encoding of the two configurations.
- *Repeat* asserts that the encoding of an accepting configuration is repeated until the maximal id is reached
- *Loop* defines the period of the word-model, which may only contain $\bot$.

All these properties can be expressed with polynomially-sized formulas, which can be found in the appendix. Furthermore, we need a formula to specify technical details: atomic propositions encoding the id's are not allowed to appear in the configurations and vice versa, symbols such as \$ and # only to appear as separators, each separator appears $p(n)$ times every d positions, configuration encodings are represented by singleton sets of letters in Σ with the exception of one set that contains a symbol from Q to determine the head position and the state of $\mathcal{M}$, etc.

For $k = l_r \cdot (l_r + 3) + 1$, each accepting run of $\mathcal{M}$ on w corresponds to exactly one k-word-model of $\varphi^w_{\mathcal{M}}$ that encodes the run in its prefix. Thus, the number of k-word-models is equal to the number of accepting runs of $\mathcal{M}$ on w. The formula $\varphi^w_{\mathcal{M}}$ can be obtained in polynomial time in $|w| + |\mathcal{M}|$, and k (thus also its unary encoding) is polynomial in $|w|$.

To show that the problem is in #P we define a nondeterministic polynomial time Turing machine $\mathcal{M}$ as follows. $\mathcal{M}$ guesses a prefix u and a period v of an ultimately periodic word $u.v^\omega$ with $|u.v| = k$, and checks deterministically in polynomial time [10], whether $u.v^\omega$ satisfies φ. Hence, for each k-word-model (u, v) of φ there is exactly one accepting run of $\mathcal{M}$. Thus, counting the k-word-

models of φ can be done by counting the accepting runs of $\mathcal{M}$ on the input (k, φ).

4.2 The Case of Binary Encodings.

Now, we consider the word counting problem for binary bounds. As the input is more compact, we have to deal with a larger complexity class.

Theorem 4. *The following problem is #*PSPACE*-complete: Given an LTL formula φ and a bound k (in binary), how many k-word-models does φ have?*

Proof. The hardness proof is similar to the one for Theorem 3: for a nondeterministic polynomial space Turing machine $\mathcal{M}$ bounded by a polynomial $p(n)$ and an input word w we can define a formula $\varphi^w_{\mathcal{M}}$ in the same way as in Theorem 3. The reason lies in that the size of configurations remains polynomial and the exponential number of configurations in a run can still be counted with a binary counter of polynomial size, i.e., we only have to use more bits b_j to encode the id's. Furthermore, we have to choose $k = 2^{p'(n)}(p(n) + 3) + 1$ which can still be encoded using polynomially many bits. Here, $p'(n)$ is a polynomial (which only depends on $\mathcal{M}$) such that $\mathcal{M}$ terminates in at most $2^{p'(n)}$ steps on inputs of length n.

For the proof of the upper bound we cannot just guess a k-model in polynomial space as in Theorem 3, since the bound k is encoded in binary. Instead, we guess and verify the model on-the-fly relying on standard techniques for LTL model checking.

Formally, we construct a nondeterministic polynomial space Turing machine $\mathcal{M}$ which guesses a k-word-model (u, v) by guessing $u\$v = w(0) \cdots w(i-1)\$$ $w(i) \cdots w(k-1)$ symbol by symbol in a backwards fashion. Here, \$ is a fresh symbol to denote the beginning of the period. To meet the space requirement, $\mathcal{M}$ only stores the currently guessed symbol $w(j)$, discards previously guessed symbols, and uses a binary counter to guess exactly k symbols.

To verify whether $u.v^\omega$ satisfies φ, $\mathcal{M}$ also creates for every j in the range $0 \leq j < k$ a set C_j of subformulas of φ with the intention of C_j containing exactly the subformulas which are satisfied in position j of $u.v^\omega$. Due to space-requirements, $\mathcal{M}$ only stores the set C_{k-1} as well as the sets C_j and C_{j+1}, if $w(j)$ is the currently guessed symbol. The set C_{k-1} is guessed by $\mathcal{M}$ and the sets C_j for $j < k-1$ are uniquely determined by the following rules:

- The membership of atomic propositions in C_j is determined by $w(j)$, i.e., $C_j \cap AP = w(j)$.
- Conjunctions, disjunctions, and negations can be checked locally for consistency, e.g., $\neg\psi \in C_j$ if and only if $\psi \notin C_j$.
- $\bigcirc$-formulas are propagated backwards using the following equivalence: $\bigcirc\psi \in C_j$ if and only if $\psi \in C_{j+1}$ (recall that $\mathcal{M}$ stores C_j and C_{j+1}).
- $\mathcal{U}$-formulas are propagated backwards using the following equivalence: $\psi_0 \mathcal{U} \psi_1 \in C_j$ if and only if $\psi_1 \in C_j$ or $\psi_0 \in C_j$ and $\psi_0 \mathcal{U} \psi_1 \in C_{j+1}$.

– $\mathcal{V}$-formulas can be rewritten into $\mathcal{U}$-formulas.

Once $\mathcal{M}$ has guessed the complete period $v = w(i) \cdots w(k-1)$ it also checks that the guess of C_{k-1} is correct (recall that C_{k-1} is not discarded), which is the case if the following two requirements are met:

– For every subformula $\bigcirc\psi$ we have $\bigcirc\psi \in C_{k-1}$ if and only if $\psi \in C_i$.
– For every subformula $\psi_0 \mathcal{U} \psi_1$ we have $\psi_0 \mathcal{U} \psi_1 \in C_{k-1}$ if and only if $\psi_1 \in C_{k-1}$ or $\psi_0 \in C_{k-1}$ and $\psi_0 \mathcal{U} \psi_1 \in C_i$. Furthermore, we have to require that $\psi_0 \mathcal{U} \psi_1 \in C_j$ for some j in the range $i \leq j < k$ implies $\psi_1 \in C_{j'}$ for some j' in the range $i \leq j' < k$. The latter condition can be checked on-the-fly while computing the C_j's.

A straightforward structural induction over the construction of φ shows that we have $\psi \in C_j$ if and only if $w(j)w(j+1) \cdots w(k-1)v^\omega \models \psi$ for every subformula ψ of φ. Hence, $u.v^\omega$ is a model of φ if and only if $\varphi \in C_0$. Thus, $\mathcal{M}$ accepts if this is the case.

5 Counting Tree-Models

In this section, we consider the tree counting problem for unary and binary bounds. There are at least doubly-exponentially many trees of height k. Hence, if k is encoded in binary, there are at least triply-exponentially many (in the size of the encoding of k) k-tree-models of a tautology. In order to capture these cardinalities using counting classes, we have to consider machines with that many runs, i.e., exponential time and exponential space machines.

In our hardness proofs, we again construct formulas $\varphi^w_{\mathcal{M}}$ that encode accepting runs of $\mathcal{M}$ on w in trees. We choose binary trees, i.e., we consider a singleton set I of input propositions. Recall that the power set of I is used to (deterministically) label the edges in the tree. In the following, we identify the two elements of 2^I with the directions `left` and `right`. Note that we have to formalize the structure of our models and have to encode the runs of the machines using LTL. The semantics require a formula to be satisfied on all paths, which requires us to write conditional formulas of the form "if the path has a certain form, then some property is satisfied". We use two types of formulas: the ones of the first type describe the structure of the tree (e.g., it is complete and the targets of the back-edges) while the ones of the second type encode the actual run relying on this structure. The formulas of type one often assign addresses to nodes (sequences of bits that uniquely identify a leaf).

In the word case, we encoded runs of Turing machines whose configurations are of polynomial length. Hence, the distance between encodings of a tape cell in two successive configurations could be covered by a polynomial number of next-operators. Here, configurations are of exponential size. Thus, the challenge is to encode a run in a tree-model such that properties of two successive configurations can still be encoded by an LTL formula of polynomial size. We present two such encodings, one for unary and one for binary bounds.

For the upper bounds we show that there are appropriate nondeterministic machines that guess a finite tree with back-edges and model check it deterministically against φ, i.e., the number of accepting runs on k and φ is equal to the number of k-tree-models of φ.

5.1 The Case of Unary Encodings

First, we consider tree-model counting for unary bounds.

Theorem 5. *The following problem is* #EXPTIME*-complete: Given an LTL formula φ and a bound k (in unary), how many k-tree-models does φ have?*

Proof. We start with the hardness proof. Let $\mathcal{M} = (Q, q_\iota, Q_F, \Sigma, \delta)$ be a one-tape nondeterministic exponential time Turing machine. Let $\mathcal{M}$ be $2^{p(n)}$-time bounded for a polynomial p and let $w = w_0 \cdots w_{n-1}$ be an input to $\mathcal{M}$. We construct an LTL formula $\varphi^w_\mathcal{M}$ and define a bound k, both polynomial in $|w|$ and $|\mathcal{M}|$, such that the number of accepting runs of $\mathcal{M}$ on w is equal to the number of k-tree-models of $\varphi^w_\mathcal{M}$.

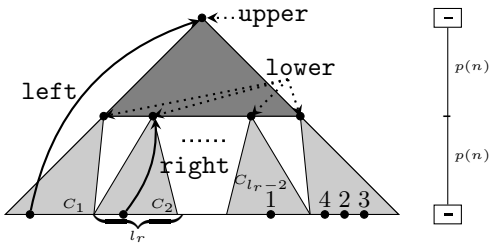

Fig. 2. Encoding an exponentially long run in a tree-model of polynomial height. The configurations are encoded in the lower-trees (light gray subtrees).

A run of $\mathcal{M}$ is encoded in the leaves of a binary tree-model. Let $l_r = 2^{p(n)}$ be the maximal length of a run of $\mathcal{M}$ on w, which also bounds the size of a configuration. We choose $k = 2p(n)$ to be the height of our tree-models. By using a formula labeling each of the first k levels of the tree by a unique proposition we enforce that every model of height k is complete. Thus, it has l_r^2 many leaves, enough to encode a run of maximal length. Figure 2 shows the structure of our tree-model.

Each configuration in the run is encoded in the leaves of a subtree of height $p(n)$, referred to as a *lower*-tree (depicted by the light gray trees). The lower-trees are uniquely determined by a leaf of the *upper*-tree (depicted in dark gray), which is the root of the lower-tree. By giving the leaves of the upper-tree id's, we also obtain unique id's for each of the lower-trees. These id's are used to enumerate the configurations of the run, i.e., two neighboring lower-trees encode two successive configurations of the run. The id's can be determined by a binary counter with polynomially many bits. We also provide each leaf in a lower-tree with a unique id within this lower-tree. This is used to compare the contents of a tape cell in two successive configurations by comparing the labels of leaves with the same leaf id in two successive lower-trees. Thus, every leaf stores the id encoding of the configuration it is part of and the number of the cell it encodes.

Recall that in a tree-model each leaf has a back-edge for every direction. For the direction `left` we require a transition to the root of the upper-tree,

and for `right` a transition to the root of the own lower-tree. This enables us to compare two leaves in a lower-tree, or two leaves with the same id in two different lower-trees, with polynomially large formulas.

The following formulas define the structure of our tree-models as explained above and also provide the nodes of the tree with correct id's. We begin with $\mathit{Addr}(\texttt{root},a_1,\ldots,a_d)$ which specifies a unique id for each leaf of a complete binary tree of height d using bits $a_1,\ldots,a_d$, and provides the root of the tree with a label `root`. The id of a node depends on the sequence of `left` and `right` edges on the path from the root to this node, which is encoded in the bits $a_1,\ldots,a_d$:

$$\mathit{Addr}(\texttt{root},a_1,\ldots,a_d) = \texttt{root} \wedge \bigwedge_{i=0}^{d-1} \big(\bigcirc^i(\texttt{left} \rightarrow \bigcirc^{d-i} \neg a_{i+1}) \wedge \bigcirc^i(\texttt{right} \rightarrow \bigcirc^{d-i} a_{i+1}) \big)$$

We use the formula $\mathit{Addr}(\texttt{upper},u_1,\ldots,u_{p(n)})$ to address the upper-tree. This gives each lower-tree a unique id via the id of its root. We also supply each node in a lower-tree with the id of its root in the upper-tree: $\bigwedge_{p(n)\leq i<k} \bigcirc^i(\bigwedge_{j=1}^{p(n)}(u_j \leftrightarrow \bigcirc u_j))$. Furthermore, we use the formula $\bigcirc^{p(n)} \mathit{Addr}(\texttt{lower},l_1,\ldots,l_{p(n)})$ to assign every leaf in a lower-tree a unique id within its lower-tree which essentially encodes the number of the tape cell it encodes. The next two formulas define the back-edges of the lower-trees. From each leaf, the `left` transition leads back to the root of the upper-tree (recall that back-edges lead from a leaf to an ancestor), i.e., $\bigcirc^k(\texttt{left} \rightarrow \bigcirc \texttt{upper})$, and the `right` transition to the root of the lower-tree, i.e., $\bigcirc^k(\texttt{right} \rightarrow \bigcirc \texttt{lower})$. After setting up the structure of the trees, it remains to show how we encode a run in the leaves. We proceed with the same scheme as in the word case, and use the formula $\Delta_h(a_1,\ldots,a_m)$ which is satisfied, if and only if the bits $a_1,\ldots,a_m$ encode the number $h < 2^m$.

- The formula *Init* encodes the initial configuration in the lower-tree with id 0.

$$\bigcirc^k \big[\Delta_0(u_1,\ldots,u_{p(n)}) \rightarrow \big((\Delta_0(l_1,\ldots,l_{p(n)}) \rightarrow q_\iota) \wedge \bigwedge_{0\leq j<n} (\Delta_j(l_1,\ldots,l_{p(n)}) \rightarrow w_j) \wedge ((\bigwedge_{0\leq j<n} \neg\Delta_j(l_1,\ldots,l_{p(n)})) \rightarrow \text{␣}) \big)\big]$$

- The formula *Accept* checks whether the rightmost lower-tree encodes an accepting configuration: $\bigcirc^k((\Delta_{l_r}(u_1,\ldots,u_{p(n)}) \wedge \bigvee_{q\in Q} q) \rightarrow \bigvee_{q\in Q_F} q)$.
- The formulas $\mathit{config}_{q,\alpha}$ and config_α for states q and symbols α encode the transition relation. For a leaf with labels q and α (leaf 1 in Figure 2) and a transition $(q,\alpha,q',\beta,\texttt{dir})$, we have to check three leaves in the next lower-tree, namely, the leaf with the same id (leaf 2) has to be labeled with β, and depending on `dir` either the successor leaf (leaf 3) or the predecessor leaf

(leaf 4) has to be labeled with q'. The premise of the following formula only holds for paths that visit these leaves in the order given above, i.e., paths that lead to a leaf in a lower-tree, loop back to the root of the full tree and then lead to the same leaf id in the successor lower-tree (this takes $k+1$ edges), loop back to the root of this lower-tree and visit the leaf to the right (this takes $p(n)+1$ edges), back to the root of this lower-tree again and then to the leaf to the left (this takes $p(n)+1$ edges). To specify such a path, we use the formula *Inc* to reach the successor leaf and a dual formula called *Dec* to reach the predecessor leaf. This formula implements a decrement of a nonzero counter. Note that we have to require the paths to visit the successor and predecessor leaf in the next lower-tree, i.e., we have to check the bits u_j to reach the next lower-tree and the bits l_j to reach the leaves. Thus, $\mathit{config}_{q,\alpha}$ for $q \in Q \setminus Q_F$ is given by:

$$\begin{aligned}\bigcirc^k \big[\ & q \wedge \alpha \wedge \mathit{Inc}(u_1,\dots,u_{p(n)},k+1) \wedge \textstyle\bigwedge_{i=1}^{p(n)} l_i \leftrightarrow \bigcirc^{k+1} l_i) \\ & \wedge \mathit{Inc}(u_1,\dots,u_{p(n)},k+p(n)+2) \wedge \mathit{Inc}(l_1,\dots,l_{p(n)},k+p(n)+2) \\ & \wedge \mathit{Inc}(u_1,\dots,u_{p(n)},k+2p(n)+3)) \wedge \mathit{Dec}(l_1,\dots,l_{p(n)},k+2p(n)+3) \\ & \rightarrow \bigvee_{(q',\beta,\ \mathtt{dir})\in\delta(q,\alpha)} (\bigcirc^{k+1}\beta \wedge \bigcirc^{(k+1)+c_{\mathtt{dir}}(p(n)+1)} q')\ \big]\end{aligned}$$

Here, we have $c_{\mathtt{dir}} = 1$, if $\mathtt{dir} = 1$, and $c_{\mathtt{dir}} = 2$, if $\mathtt{dir} = -1$.
The formula config_α determines the relation between the other tape cells' contents, namely where the head is not pointing to:

$$\begin{aligned}\bigcirc^k\Big(\bigvee_{i=1}^{p(n)} \neg u_i \wedge \Big(\bigwedge_{q\in Q\setminus Q_F} \neg q\Big) \wedge \alpha \wedge \mathit{Inc}(u_1,\dots,u_{p(n)},k+1) \\ \wedge\Big(\bigwedge_{i=1}^{p(n)} l_i \leftrightarrow \bigcirc^{k+1} l_i\Big) \rightarrow \bigcirc^{k+1}\alpha\Big)\end{aligned}$$

– The formula *Repeat* repeats accepting states in the next lower-tree, if the id of the current lower-tree is not maximal. The repetition of the letters is being taken care of by config_α.

$$\begin{aligned}\bigcirc^k\Big[\Big(\bigvee_{i=1}^{p(n)} \neg u_i \wedge \mathit{Inc}(u_1,\dots,u_{p(n)},k+1) \wedge \bigwedge_{i=1}^{p(n)} (l_i \leftrightarrow \bigcirc^{k+1} l_i)\Big) \rightarrow \\ \Big(\bigwedge_{q_f\in Q_F} q_f \rightarrow \bigcirc^{k+1} q_f\Big)\Big]\end{aligned}$$

Similar to the word case we need some additional formulas to prevent atomic propositions of configurations to appear elsewhere in the tree to guarantee the one-to-one relation between runs and tree-models. For example to prevent a state label from appearing twice in a configuration we use a formula that asserts that from a leaf in which a state is encoded, no other leaf with a state label is

reachable within $p(n) + 1$ steps, i.e., in the same lower-tree. This ensures that every configuration has exactly one state.

To show that the problem is in #EXPTIME we define a nondeterministic exponential time Turing machine $\mathcal{M}$ as follows. $\mathcal{M}$ guesses a tree of height k (which is of exponential size) and checks whether it satisfies φ using the classical model checking algorithm: $\mathcal{M}$ constructs the Büchi automaton recognizing the language of $\neg\varphi$ and checks whether the product of the tree and the automaton has an empty language. The automaton and the product are of exponential size and the emptiness check can be performed in deterministic polynomial time (in the size of the product). Hence, $\mathcal{M}$ runs in exponential time in k and the size of φ. For each k-tree-model of φ, there is exactly one accepting run in $\mathcal{M}$. Thus, counting the k-tree-models of φ can be done by counting the accepting runs of $\mathcal{M}$ on the input (k, φ). □

5.2 The Case of Binary Encodings.

In this section, we consider tree-model counting for binary bounds. Since the bound is encoded compactly, the trees we work with have exponential height and therefore doubly-exponential size. Unfortunately, our upper and lower bounds do not match (see the discussion in the next section).

Theorem 6. *The following problem is #EXPSPACE-hard and in #2EXPTIME: Given an LTL formula φ and a bound k (in binary), how many k-tree-models does φ have?*

Proof. Let $\mathcal{M} = (Q, q_\iota, Q_F, \Sigma, \delta)$ be a one-tape nondeterministic exponential space Turing machine and let $w = w_0 \cdots w_{n-1}$ be an input to $\mathcal{M}$. Furthermore, let $l_c = 2^{p(n)} - 2$ be the maximal configuration length (for some polynomial p) and let $l_r = 2^{2^{p'(n)}}$ be the maximal length of a run of $\mathcal{M}$ on w (p' is a polynomial which only depends on $\mathcal{M}$).

We choose $k = m \cdot 2^{p'(n)}$ to be the height of our tree-models, where m is the smallest power of two greater than $p(n)$. Figure 3 shows the main structure of our tree-models. We use nonbalanced binary trees that are composed of trees of height m. We refer to the latter trees as the *inner*-trees. The outermost leaves of an inner-tree are inner nodes and the others are leaves in the tree-model. Hence, each inner-tree has two children, which are again inner-trees rooted at the leftmost respectively the rightmost leaf.

In each inner-tree, we will encode a configuration in a similar way as in the unary case (Theorem 5), namely in the leaves (except the two leaves serving as roots for other inner trees, which explains the -2 in the definition of l_c). We encode the configurations of a run in the tree-model such that we traverse the inner-trees in a depth-first search manner (DFS). In Figure 3, we can see how a run of 16 configurations can be encoded in a tree-model with four layers of inner-trees. To encode the DFS structure, we label each root of an inner-tree with its *level* (the number of inner-tree ancestors) and with its so-called *right-child-depth*: the number of right-child-inner-trees visited since the last left child

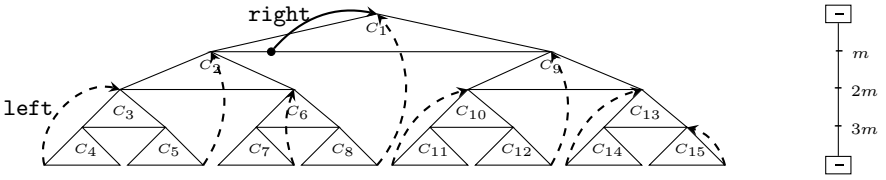

Fig. 3. Tree-model with DFS structure.

to reach this tree (e.g., this value is 0 for the left children C_1, C_2, C_3, C_7; it is 1 for C_6 and 3 for C_{15}). This will help to determine the next inner-tree in line in the DFS structure. We need a polynomial number of bits to encode these addresses. With the `right` transition we allow the leaves of an inner-tree to reach its root and we use `left` in the inner-tree of maximal level to reach the parent of the next inner-tree in DFS order. In this way, the distance between the encoding of a tape cell in two successive configurations is polynomial.

As the distance between an inner-tree and its successor is polynomial, the formulas for encoding the run in the tree-model adapt the ideas of the formulas in the unary case with slight modifications that deal with the DFS order of inner-trees. A detailed description of the construction can be found in the appendix.

The upper bound is proved using the same algorithm as in the proof of Theorem 5.

6 Discussion

We investigated the complexity of the model counting problem for specifications in Linear-time temporal logic. The word-model counting problems are #P-complete (for unary bounds) respectively #PSPACE-complete (for binary bounds) while the tree-model counting problems are #EXPTIME-complete respectively #EXPSPACE-hard and in #2EXPTIME, i.e., the exact complexity of the tree-model counting problem for binary bounds is open.

The problem we face trying to lower the upper bound is that we cannot guess the complete tree-model in nondeterministic exponential space. To meet the space-requirements, we have to construct it step by step, as in the proof of the corresponding upper bound in the word case. However, the correctness of the on-the-fly model checking procedure described there relies on the fact that the model is an ultimately-periodic word. It is open whether the technique can be extended to tree-models. On the other hand, if we try to raise the lower bound, we have to encode doubly-exponential time Turing machines, which seems challenging using polynomially-sized LTL formulas.

To conclude, let us mention another variation of the model counting problem: counting arbitrary transition systems, where the bound k now refers to the size of the transition system. For unary bounds, the problem is #P-hard, which can be shown by strengthening Theorem 3, and in $\#_o$PSPACE, since LTL model checking is in PSPACE. For binary bounds, the construction presented in Theorem 5 yields

#EXPTIME-hardness and the problem is in #EXPTIME, which can be shown by adapting the algorithm presented in the theorem.

Acknowledgments.

We would like to thank Markus Lohrey and an anonymous reviewer for bringing Ladner's work on polynomial space counting [11] to our attention.

References

1. Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, New York, NY, USA, 1st edition, 2009.
2. Armin Biere. Bounded model checking. In *Handbook of Satisfiability*, pages 457–481. IOS Press, 2009.
3. Roderick Bloem, Hans-Jürgen Gamauf, Georg Hofferek, Bettina Könighofer, and Robert Könighofer. Synthesizing robust systems with RATSY. In Doron Peled and Sven Schewe, editors, *SYNT*, volume 84 of *EPTCS*, pages 47–53. Open Publishing Association, 2012.
4. Aaron Bohy, Véronique Bruyère, Emmanuel Filiot, Naiyong Jin, and Jean-François Raskin. Acacia+, a tool for LTL synthesis. In P. Madhusudan and Sanjit A. Seshia, editors, *CAV*, volume 7358 of *LNCS*, pages 652–657. Springer, 2012.
5. Jerry R. Burch, Edmund M. Clarke, Kenneth L. McMillan, David L. Dill, and L. J. Hwang. Symbolic model checking: 10^{20} states and beyond. *Inf. Comput.*, 98(2):142–170, 1992.
6. Rüdiger Ehlers. Unbeast: Symbolic bounded synthesis. In *TACAS*, volume 6605 of *LNCS*, pages 272–275. Springer-Verlag, 2011.
7. Bernd Finkbeiner and Sven Schewe. Bounded synthesis. *International Journal on Software Tools for Technology Transfer*, 15(5-6):519–539, 2013.
8. Bernd Finkbeiner and Hazem Torfah. Counting models of linear-time temporal logic. In Adrian Horia Dediu, Carlos Martín-Vide, José Luis Sierra-Rodríguez, and Bianca Truthe, editors, *LATA*, volume 8370 of *LNCS*, pages 360–371. Springer, 2014.
9. Lane A. Hemaspaandra and Heribert Vollmer. The satanic notations: counting classes beyond #P and other definitional adventures. *SIGACT News*, 26(1):2–13, 1995.
10. Lars Kuhtz and Bernd Finkbeiner. LTL path checking is efficiently parallelizable. In Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris Nikoletseas, and Wolfgang Thomas, editors, *ICALP*, volume 5556 of *LNCS*, pages 235–246. Springer, 2009.
11. Richard E. Ladner. Polynomial space counting problems. *SIAM J. Comput.*, 18(6):1087–1097, 1989.
12. Maciej Liśkiewicz, Mitsunori Ogihara, and Seinosuke Toda. The complexity of counting self-avoiding walks in subgraphs of two-dimensional grids and hypercubes. *Theor. Comput. Sci.*, 1-3(304):129–156, 2003.
13. Michael L. Littman, Stephen M. Majercik, and Toniann Pitassi. Stochastic boolean satisfiability. *Journal of Automated Reasoning*, 27:2001, 2000.
14. Markus Lohrey and Manfred Schmidt-Schauß. Processing succinct matrices and vectors. In Edward A. Hirsch, Sergei O. Kuznetsov, Jean-Éric Pin, and Nikolay K. Vereshchagin, editors, *CSR*, volume 8476 of *LNCS*, pages 245–258. Springer, 2014.

15. Daniel Morwood and Daniel Bryce. Evaluating temporal plans in incomplete domains. In Jörg Hoffmann and Bart Selman, editors, *AAAI*. AAAI Press, 2012.
16. Amir Pnueli. The temporal logic of programs. In *STOC*, SFCS '77, pages 46–57, Washington, DC, USA, 1977. IEEE Computer Society.
17. A. Prasad Sistla and Edmund M. Clarke. The complexity of propositional linear temporal logics. *J. ACM*, 32(3):733–749, 1985.
18. A.Prasad Sistla. Safety, liveness and fairness in temporal logic. *Formal Aspects of Computing*, 6(5):495–511, 1994.
19. Leslie G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8:189–201, 1979.

A Appendix

A.1 Proof of Theorem 3

In this section, we present the formulas omitted in the proof of #P-hardness of the word counting problem for unary bounds.

We start with the formula *Id*, which uses the formula *Inc* that enforces an increment of a binary counter. For later reuse, *Inc* is parameterized by the propositions $b_1, \ldots, b_\ell$ encoding the bits and the distance d between the two positions to be compared. Intuitively, the different subformulas distinguish whether the increment ripples through to the current bit b_i or not. Note that the increment property only has to hold if there is no overflow of the counter.

$$\begin{aligned} \mathit{Inc}(b_1,\ldots,b_\ell,d) = (\bigvee_{0<i\leq\ell} \neg b_i) \rightarrow \bigwedge_{0<i\leq\ell} \big[&((\neg b_i \wedge \bigwedge_{i<j\leq\ell} b_j) \rightarrow \bigcirc^d \; b_i) \\ &\wedge((\neg b_i \wedge \neg \bigwedge_{i<j\leq\ell} b_j) \rightarrow \bigcirc^d \neg b_i) \\ &\wedge((\; b_i \wedge \bigwedge_{i<j\leq\ell} b_j) \rightarrow \bigcirc^d \neg b_i) \\ &\wedge((\; b_i \wedge \neg \bigwedge_{i<j\leq\ell} b_j) \rightarrow \bigcirc^d \; b_i) \big] \end{aligned}$$

Now, the formula *Id* is defined by initializing the counter to zero and always requiring an increment after a \$-separator:

$$\mathit{Id} = \$ \;\wedge\; \bigcirc(\bigwedge_{0<j\leq l_c} \neg b_j) \;\wedge\; \Box(\$ \rightarrow \bigcirc \mathit{Inc}(b_1,\ldots,b_{l_c},d))$$

We continue with the formula *Init*. In the initial configuration the tape of $\mathcal{M}$ contains the input word w, the head is on the first cell, and $\mathcal{M}$ is in its initial state:

$$\mathit{Init} = \bigcirc^2(\# \;\wedge\; \bigcirc q_\iota \;\wedge\; (\bigwedge_{0\leq j<n} \bigcirc^j w_j) \;\wedge\; (\bigwedge_{n\leq j\leq l_r} \bigcirc^j \text{␣}))$$

The symbol ␣ refers to the blank cells of the tape.

The formula *Accept* considers the maximal id and checks whether it is followed by an accepting configuration:

$$\mathit{Accept} = \Box(\$ \wedge \bigcirc(\bigwedge_{0<j\leq l_c} b_j) \;\rightarrow\; \bigvee_{q\in Q_F} \bigvee_{0<j\leq l_r} \bigcirc^{j+2} q)$$

For atomic propositions $q \in Q \setminus Q_F$ and $\alpha \in \Sigma$ a formula $\mathit{config}_{q,\alpha}$ asserts the relation between the states, the head positions, and the content of the cell where the head is pointing to in two successive configurations:

$$\mathit{config}_{q,\alpha} = \Box(\; q \wedge \alpha \rightarrow \bigvee_{(q',\beta,\mathrm{dir})\in\delta(q,\alpha)} \bigcirc^d \beta \wedge \bigcirc^{d+\mathrm{dir}} q' \;)$$

Another formula config_α asserts the relation of the other tape cells of successive configurations; the content of these cells is copied, unless the id is maximal[3]:

$$\mathit{config}_\alpha = \square(\$ \wedge \bigcirc(\bigvee_{0<j\leq l_c} \neg b_j) \rightarrow \bigwedge_{0<j\leq l_r} \bigcirc^{j+2}((\bigwedge_{q\in Q\setminus Q_F} \neg q) \wedge \alpha \rightarrow \bigcirc^d \alpha))$$

Config is the conjunction of all formulas config_α and $\mathit{config}_{q,\alpha}$.

The formula *Repeat* requires an accepting configuration to be repeated if the id is not yet maximal. The repetition of the letters is taken care of by the formulas config_α. Hence, *Repeat* only requires to copy the state and the head position.

$$\mathit{Repeat} = \square(\$ \wedge (\bigcirc \bigvee_{0<j\leq l_c} \neg b_j) \rightarrow \bigwedge_{q_f\in Q_F} \bigwedge_{0<j\leq l_r} \bigcirc^{j+2}(q_f \rightarrow \bigcirc^d q_f))$$

Finally, the period of the model is fixed by the formula *Loop* which requires the symbol $\perp$ to be repeated after reaching the configuration with the maximal id:

$$\mathit{Loop} = \square(\$ \wedge \bigcirc(\bigwedge_{0<j\leq l_c} b_j) \rightarrow \bigcirc^{l_r+2} \ \square\perp)$$

A.2 Proof of Theorem 6

The following formulas define the structure of our tree models and also provide them with the correct level and right-child-depth. We use propositions $\iota_1, \ldots, \iota_m$ to give *i*ds to the leaves of an inner-tree. For the $2^{p'(n)}$ different levels of inner-trees in our tree models we use propositions $l_1, \ldots, l_h$, where $h = p'(n)$, to encode for each inner-tree its *l*evel. We also give internally for each node in an inner-tree its *d*epth in this tree via $d_1, \ldots, d_{\log(m)}$ (remember that m is a power of 2). Propositions $r_1, \ldots, r_h$ are used to determine the inner-tree's *r*ight-child-depth. The maximum right-child-depth that can be reached is $2^{p'(n)}$, namely for the rightmost inner-tree at the maximal level.

We start by labeling each root of an inner-tree (and no other vertices) with the label `root`:

$$\texttt{root} \wedge$$
$$\square\big[\ \big(\texttt{root} \wedge \neg(\bigwedge_{0<i\leq h} l_i) \wedge ((\bigwedge_{0\leq j<m} \bigcirc^j \texttt{left}) \vee (\bigwedge_{0\leq j<m} \bigcirc^j \texttt{right}))\big) \leftrightarrow \bigcirc^m \texttt{root}\big)\ \big]$$

The negation of the big conjunction over propositions l_i is used to exclude inner-trees of the last level. How the levels are defined in the tree is shown in more detail in the formula *depth*.

[3] Note that this is not necessary for $\mathit{config}_{q,\alpha}$ since the machine terminates in at most $p(n)$ steps

We encode a configuration in the leaves of an inner-tree in the same way as in the unary case (Lemma 5). Therefore, we provide the leaves with unique id's, which enable us to compare the contents of tape cells in two successive configurations by comparing leaves with the same id in two successive inner-trees. To this end, we use the formula *Addr*, as defined in the unary case, to equip the leaves with unique id's within their inner-tree:

$$\Box(\texttt{root} \rightarrow \mathit{Addr}(\texttt{root}, \iota_1, \ldots, \iota_m))$$

The next formula uses the propositions $l_1, \ldots, l_h$ to supply each inner-tree with its level, which is equal to the number of `root` labels visited from the root to this inner-tree. Line (1) assigns the root of the tree model with level 0. Line (2) assigns each node in an inner-tree, with the exception of the leaves labeled with `root` (the outermost leaves in all levels but the last), with the level of its inner-tree root. Line (3) gives, inductively, each root of an inner-tree its inner-tree level.

$$\begin{aligned}
\mathit{depth} = & \bigwedge_{0<j\leq h} \neg l_j && (1)\\
& \wedge \Box(\texttt{root} \rightarrow \bigwedge_{0\leq j<m} \bigcirc^j(\bigcirc\neg\texttt{root} \rightarrow \bigwedge_{0<j\leq h} (l_j \leftrightarrow \bigcirc l_j)) && (2)\\
& \wedge \Box(\texttt{root} \wedge (\bigvee_{0<j\leq h} \neg l_j) \wedge \bigcirc^m \texttt{root} \rightarrow \mathit{Inc}(l_1, \ldots, l_h, m) && (3)
\end{aligned}$$

If *depth* is satisfied then the tree contains a doubly-exponential number of inner-trees, enough to encode a run of $\mathcal{M}$ on w.

The following formula gives each inner-tree its right-child-depth in the tree model. If we are at a root of an inner-tree we reach the root of its left and right child in m steps via the outermost leaves. The formula counts the number of right-children visited along the way including the visited tree (Line (3)). Every time we visit a left-child the counter is reset (Line (4)). We again supply each node in an inner-tree, with the exception of the outermost leaves (Line (5)), with the right-child-depth of this tree. However the outermost leaves of all maximal-level inner-trees are labeled with the right-child-depth of the inner-tree.

$$\begin{aligned}
\mathit{rLevel} = & \bigwedge_{0<j\leq h} \neg r_j && (1)\\
& \wedge \Box(\texttt{root} \wedge \bigvee_{0<j\leq h} \neg l_j \rightarrow && (2)\\
& \qquad ((\bigwedge_{0<j\leq m} \bigcirc^j \texttt{right}) \rightarrow \mathit{Inc}(r_1, \ldots, r_h, m)) && (3)\\
& \qquad \wedge ((\bigwedge_{0<j\leq m} \bigcirc^j \texttt{left}) \rightarrow \bigwedge_{0<j\leq h} \neg r_j)) && (4)\\
& \wedge \Box(\texttt{root} \rightarrow \bigwedge_{0\leq j<m} \bigcirc^j(\bigcirc\neg\texttt{root} \rightarrow \bigwedge_{0<j\leq h} r_j \leftrightarrow \bigcirc r_j))) && (5)
\end{aligned}$$

Now that we have defined all the id's we need we show how to route the transitions at the leaves to the DFS positions as described earlier. To compute the correct node we only need to compute its level, because in the definition of our tree models we force a back-edge to go to an ancestor node. This level can be computed as follows. If we are at a leaf of an inner-tree in the maximal level, reached by j many right children since the last left child the back-edge goes to the root with level $2^{p'(n)} - (j+1)$ (cf. Figure 3). This can be formulated by incrementing the right-child-depth described with bits $r_1, \ldots, r_h$ and inverting the bits of the result. The propositions $d_1, \ldots, d_{\log(m)}$ are used to talk about the leaves of an inner-tree. If we are at a leaf in an inner-tree of maximal level we move with direction `left` to the next DFS position, namely the root of the next inner-tree in the DFS Line (computed by *addNflip*):

$$\mathit{DFS} =$$
$$\square\Big(\bigwedge_{0<j\leq\log(m)} d_j \wedge \bigwedge_{0<j\leq h} l_j \wedge \texttt{left} \rightarrow \bigcirc\texttt{root} \wedge \mathit{addNflip}(l_1, \ldots, l_h, r_1, \ldots, r_h)\Big)$$

where:

$$\begin{aligned}\mathit{addNflip}(l_1, \ldots, l_c, r_1, \ldots, r_c) = \bigwedge_{0<i\leq c} \Big[&\Big((\neg r_i \wedge \bigwedge_{i<j\leq c} r_j) \rightarrow \bigcirc\neg l_i\Big)\\ &\wedge\Big((\neg r_i \wedge \neg \bigwedge_{i<j\leq c} r_j) \rightarrow \bigcirc\ \ l_i\Big)\\ &\wedge\Big((\ \ r_i \wedge \bigwedge_{i<j\leq c} r_j) \rightarrow \bigcirc\ \ l_i\Big)\\ &\wedge\Big((\ \ r_i \wedge \neg \bigwedge_{i<j\leq c} r_j) \rightarrow \bigcirc\neg l_i\Big)\end{aligned}$$

Notice that *addNflip* is similar to *Inc* with the difference of flipping the bit.

Finally, a formula that asserts that from leaves of maximal level inner-trees a `right` transition goes to the root of these trees (Line (1)). For leaves of inner-trees of non-maximal level (the outermost leaves are not considered leaves of the tree model as their `right` and `left` transitions lead to subtrees), we move with both `left` and `right` to the root of their inner-trees (Line (2)):

$$\square\Big(\bigwedge_{0<j\leq\log(m)} d_j \wedge \bigwedge_{0<j\leq h} l_j \wedge \texttt{right} \rightarrow \bigcirc(\texttt{root} \wedge \bigwedge_{0<j\leq h} l_j)\Big) \tag{1}$$
$$\wedge\,\square\Big(\bigwedge_{0<j\leq\log(m)} d_j \wedge \Big(\bigvee_{0<j\leq h} \neg l_j\Big) \wedge \neg\texttt{root} \rightarrow \bigcirc\texttt{root} \wedge \Big(\bigwedge_{0<j\leq h} l_j \leftrightarrow \bigcirc l_j\Big)\Big) \tag{2}$$

Now, we present the formulas that encode the run of the Turing machine. Here, we again use the formula $\Delta_h(a_1, \ldots, a_m)$ which is satisfied, if and only if the bits $a_1, \ldots, a_m$ encode the number $h < 2^m$.

– The formula *Init*:

$$\begin{aligned} \mathit{Init} = \bigcirc^m((\Delta_1(\iota_1,\ldots,\iota_m) \rightarrow q_\iota) \\ \wedge (\bigwedge_{0\leq j<n} (\Delta_{j+1}(\iota_1,\ldots,\iota_m) \rightarrow w_j)) \\ \wedge ((\bigwedge_{0<j\leq n} (\neg\Delta_j(\iota_1,\ldots,\iota_m)) \rightarrow \text{␣}))) \end{aligned}$$

Note that we encode the input word $w = w_0 \cdots w_{n-1}$ in the leaves with id's 1 to $n-1$. In our encoding the outermost leaves of the inner-trees will not encode any tape content of the Turing machine. This is due to the fact that in some inner-trees these leaves have no back-edges from which we can directly reach the leaves of the next inner-tree in DFS order. Thus, the content in the next configuration is not accessible.

– The formula *Accept* checks the last inner-tree in DFS order for an accepting configuration. The first release formula selects the outermost right path and stops at the root of the last inner-tree. If we arrive at this root we assert that the state in the configuration of this inner-tree must be accepting:

$$(\texttt{right} \wedge \neg\Delta_{2^h-1}(l_1,\ldots,l_h))\mathcal{V}(\Delta_{2^h-1}(l_1,\ldots,l_h) \tag{1}$$

$$\rightarrow (\texttt{right} \wedge \neg\Delta_{2^h-1}(l_1,\ldots,l_h))\mathcal{V} \tag{2}$$

$$(\Delta_{2^h-1}(l_1,\ldots,l_h) \wedge \bigcirc^m(\bigvee_{q\in Q} q \rightarrow \bigvee_{q\in Q_F} q))) \tag{3}$$

– We define the formula *Next* to determine the successor inner-tree of an inner-tree, i.e., the formula holds at a vertex on a path if after m steps on this path the root of the next inner-tree in DFS order is reached:

$$\mathit{Next} = [(\bigvee_{0<j\leq m} \neg l_j) \tag{1}$$

$$\rightarrow \texttt{right} \wedge (\bigwedge_{0<j\leq m} \bigcirc^j \texttt{left})] \tag{2}$$

$$\wedge [(\bigwedge_{0<j\leq m} l_j) \tag{3}$$

$$\rightarrow \texttt{left} \wedge (\bigwedge_{0<j\leq m} \bigcirc^j \texttt{right})] \tag{4}$$

We distinguish two cases asserted by Lines (1) and (3), namely the case of an inner-tree in the maximal level and one in a non-maximal level. In the second case, the successor tree is the left child of the current inner-tree. Here, we go up to the root of the inner-tree and traverse down the left side to the root of the left child (Line (2)). If we are in the maximal level the successor tree is reached via the DFS order, which means, going to the inner-tree in DFS order and traversing down the right side to the right child (Line (4)).

– If we are at a leaf with symbol $\alpha \in \Sigma$ and state $q \in Q \setminus Q_F$ we move to the root of the next inner-tree in DFS order (this is determined by the formula *Next*). In this tree we have to check whether α is rewritten to the correct symbol and whether the head moved to the correct position. This is checked in the same way as in the proof of the unary case, namely, in three phases: we consider a path that leads to a leaf in successor inner-tree with the same leaf id, loops back up to the root of the inner-tree, and leads down to the same tree and visits the leaf to the right, loops back up and down the same tree and visits the leaf on the left (Lines (2),(3)). In Line (4), the distance $2m+1$ results from going one step in the inner-tree and then going $2m$ steps down to the leaves of the successor tree. The distance $3m+2$ results from looping in the same tree a second time, and $4m+3$ from looping a third time.
Here, we again use the formulas *Inc* and *Dec* introduced in the previous section of the appendix.

$$\begin{aligned}
\mathit{config}_{q,\alpha} = \square(q \wedge \alpha \wedge \mathit{Next} \wedge \quad & (\bigwedge_{j=1}^{m} \iota_j \leftrightarrow \bigcirc^{4m+1} \iota_j) && (1)\\
& \wedge \mathit{Inc}(\iota_1, \ldots, \iota_m, 3m+2)) \wedge \bigcirc^{2m+1}(\texttt{right} \wedge \bigcirc \texttt{root}) && (2)\\
& \wedge \mathit{Dec}(\iota_1, \ldots, \iota_m, 4m+3) \wedge \bigcirc^{3m+2}(\texttt{right} \wedge \bigcirc \texttt{root}) && (3)\\
& \rightarrow (\bigvee_{(q', \beta,\ \mathtt{dir}) \in \delta(q,\alpha)} \bigcirc^{2m+1} \beta \wedge \bigcirc^{c_{\mathtt{dir}}(m+1)-1} q') && (4)
\end{aligned}$$

where $c = 3$ for $\texttt{dir} = 1$ and $c = 4$ for $\texttt{dir} = -1$.
– The formula config_α is similar :

$$\mathit{config}_\alpha = \square((\bigwedge_{q \in Q \setminus Q_F} \neg q) \wedge \alpha \wedge \mathit{Next} \wedge (\bigwedge_{j=1}^{m} \iota_j \leftrightarrow \bigcirc^{2m+1} \iota_j) \rightarrow \bigcirc^{2m+1} \alpha)$$

Note that due to the DFS structure, we do not need to check for the configuration having a non-maximal id as in the unary case.
– Finally the formula *Repeat* propagates all final states q to the successor tree:

$$\begin{aligned}
\mathit{Repeat} = \square((\bigvee_{j=1}^{h} \neg r_j) \wedge \mathit{Next} \wedge (\bigwedge_{j=1}^{m} \iota_j \leftrightarrow \bigcirc^{2m+1} \iota_j)\\
\rightarrow (\bigvee_{q \in Q_F} (q \rightarrow \bigcirc^{2m+1} q))
\end{aligned}$$

Again we need an additional formula that for example forces the atomic propositions to appear only in the designated node or to have only one state label in each inner-tree, and some more for other technical properties. □